

# Index

## A

Abstractions in a cloud environment, 13  
Acceleration Brick Architecture (dACCEL-BRICK)  
    hybrid optical and electrical switching, 40  
    static infrastructure hosts  
        communication with remote dCOM-PUBRICKs, 43  
        dynamic hardware reconfiguration, 43  
        interfacing with the hardware accelerator, 43  
Accelerator Functional Unit (AFU), 68  
Algorithmic decomposition, 134, 135  
*Allcache* pin tool, 51  
All-programmable MPSoCs  
    typical machine learning applications, 4  
    logistic regression and K-Means, 4, 88, 89  
Amazon Web Services, 31  
Apache Hadoop, 67, 91  
Apache Spark, 67, 88, 90  
API Microversions framework, 157  
Application Lifecycle Deployment Engine (ALDE), 139–141  
Application Programming Interface (API), 92, 191, 192, 194  
AP SoC and the DRAM, 104  
ARMA-based framework, 61  
ARM Cortex-A9 MPCore processor, 68  
Automotive embedded systems for Driver Assistance (ADAS), 5, 181, 185  
Autoregressive (AR) model, 61  
Autoregressive Fractionally Integrated Moving Average (ARFIMA), 61

Autoregressive Integrated Moving Average (ARIMA), 61  
Average DC utilization, 65, 229, 231, 235, 236  
AXI-compatible port interface  
    AXI DDR controller, 43  
    PL DDR memory, 43  
AXI DMA module, 154

## B

Barcelona Supercomputing Center (BSC), 136  
Basic Linear Algebra Subprograms (BLAS), 93, 124  
Best energy-efficient node, 173  
Binary logistic regression model, 94, 95  
Bin Packing Problem (BPP), 60  
Bin-packing with overcommitment, 174  
Board Management Communication (BMC), 44  
BRAM limitations, 98  
Breeze library, 92  
Business Intelligence (BI), 92  
Business-Level Objectives (BLOs), 133

## C

Checkpoint-Restore in Userspace (CRIU)  
    application snapshots, 77  
    remote application analysis, 77  
    remote debugging, 77  
Chip Power Model (CPM), 258  
Cinder and Ceph file system, 82  
Cloud and edge computing paradigm, 129, 241  
Cloud-based IoT platforms

- Machine-to-machine communication protocols
    - IoT-specific ones, 64
    - REST, 64
    - Web-Sockets, 64
  - Cloud computing platform
    - accelerators, 38, 88, 90
    - cloudification of Internet of Things (IoT), 63
    - cloudlets, 63
    - deployment of hardware accelerators, 88
    - fog computing, 63
    - GPU, 62
  - Cloud infrastructure perimeter, 63
  - Cloudlets, 63
  - CloudLightning architecture, 5, 165–167, 179
  - Cloud orchestration software, 60
  - Cloud service model
    - Infrastructure-as-a-Service (IaaS), 62, 115
    - Platform-as-a-Service (PaaS), 62
    - Software-as-a-Service (SaaS), 62
  - CloudSuite
    - data caching benchmark, 52
  - CloudSuite benchmark suites, 3
  - Coarse-grain tasks, 137, 138
  - Coherent Accelerator Processor Interface (CAPI), 59, 68, 90
  - Coherent Attached Processor Proxy (CAPP), 68
  - Compilation frameworks
    - LegUp, 61
    - OpenCL, 61
    - ROCCC, 61
  - COMP superscalar (COMPSs), 137
  - Compute brick
    - Linux kernel with KVM modules, 49
    - state-of-the-art systems, 49
    - virtual machine management capabilities, 49
  - Compute-brick kernel, 50
  - Constrained application protocol, 64
  - Constraint Satisfaction Problem (CSP), 60
  - Container migration, 76, 77
  - Contiguous memory allocate, 99
  - Contiguous Memory Free (CMF), 100
  - Conventional ARM-based system, 204
  - Conventional low-power variation-aware design, 4, 242
  - Convolutional Neural Networks (CNNs), 80
  - CPU technologies, 1, 87
  - Cryptographic primitives
    - OpenCL implementations, 121
    - CUDA kernel, 139
    - Cumulative Distribution Function (CDF), 263
    - Cyber-Physical Systems (CPS), 3, 57, 80
- ## D
- DAG scheduler, 91
  - Data center agnostic features
    - iWARP, 118
    - RoCE, 118
    - VxLAN, 118
  - Data Center (DC) servers, 58
  - Data center disaggregation, 3, 37
  - Data Center-in-a-Box, 59
  - Data centres Optimization for energy-efficient and environmentalLy Friendly iNternet (DOLFIN), 218, 220–223, 226–228, 230, 231, 233, 235–237
  - dBOXs dBOSM optical switch, 44
  - dBRICK architecture, 40
  - DC federation, 237
  - DC Operator (DCO) hypervisor manager, 6, 221
  - DDR4 SO-DIMMs, 118
  - Dennard’s scaling, 1, 2, 87, 88
  - Device emulator, 133
  - Device Supervisor (DS), 140
  - Disaggregated memory simulator
    - DRAMSim2, 51, 52
    - Intel’s PIN framework, 51
  - disaggregated Recursive Data center in a Box (dReDBox), 3, 36, 38–42, 46, 47, 50
  - Distributed Resource Scheduler (DRS) tool, 60
  - DMTF Redfish API, 121
  - DNA sequencing application, 89
  - DRAM memory, 103, 203
  - dReDBox hardware architecture, 3, 39
  - dReDBox rack architecture
    - Compute Brick Architecture (dCOM-PUBRICK)
      - MPSoC integrates, 41
      - off-chip memory (DDR4), 41
    - Memory Brick Architecture (dMEM-BRICK)
      - Xilinx Zynq Ultrascale, 42
  - DTLB load page walks, 79
  - dTRAY architecture, 44
  - Dynamic allocation strategy, 74

- Dynamic Partial Reconfiguration (DPR), 154, 190
  - Dynamic Random-Access Memory (DRAM), 78, 104, 152, 205, 210, 262
  - Dynamic Voltage And Frequency Scaling (DVFS), 121
  - DyRACT implementation
    - Virtex 6 and Virtex 7 FPGAs, 38
- E**
- Earliest Deadline First (EDF), 191
  - ECOSCALE solution, 199
  - ECRAE policies, 73
  - Efficient Cloud Resources Allocation Engine (ECRAE), 70, 71, 73–75
  - Energy efficiency, 2–6, 38, 58–61, 64, 65, 70, 79–81, 88, 89, 104, 106, 111–114, 117, 119, 124, 126, 130–133, 147, 164, 165, 168, 169, 171, 173, 174, 182, 188, 190, 191, 193, 216, 217, 223, 228, 231–236, 241, 243, 269
  - Energy Modeller (EM), 142
  - Enhanced transmission selection of IEEE, 27
  - EU-funded research efforts
    - ECOSCALE, 38
    - hybrid programming environment
      - MPI and OpenCL, 38
    - Vineyard, 38
  - EULAG, 124, 125
  - Expected energy consumption, 232, 236
  - Expected revenue change, 230, 233, 235, 236
  - Extended Operating Points (EOP), 245
- F**
- Field-Programmable Gate Arrays (FPGAs)
    - high-performance reconfigurable computing group, 10
    - in data centre, 3, 14, 17, 23
  - Floating point Operations Per Second (FLOPS), 170
  - FPGA-based accelerators, 2, 38
  - FPGA-based hardware resources, 90
  - FPGA-based MPSoCs, 88
  - FPGA Fabric, 11, 119, 126
  - FPGA hypervisor
    - basic I/O abstraction, 16
    - Hardware Abstraction Layer (HAL), 16
    - off-chip DRAM, 16
  - FPGA SoC-based programmable cloud platform, 153
  - FPGA VNF service chain scheduler, 23
  - FPGA-based accelerators, 120
  - FPGA/GPGPU hardware, 122
- G**
- GA-generated voltage noise virus, 261
  - Genetic Algorithms (GA), 60
  - Gigabit Ethernet interface, 157
  - Google’s tensor processing unit, 150
  - Gradient descent algorithm, 98
  - Gradient Descent (GD), 95
  - Gradient descent iteration, 102
  - Graph computations (GraphX), 91
  - Graphics Processing Units (GPUs), 89, 126, 134, 140, 141, 145, 150, 165, 182, 204
  - Graphics processing units (GPUs), 111, 114, 115, 121
  - GraphX, Spark API (Application Programming Interface), 91
  - GUPS benchmark, 79
- H**
- Hard Processor System (HPS), 68
  - Hardware accelerators
    - data centres
      - cloud computing, 2
      - computer architectures, 2
      - data center design, 2
      - high performance computing, 2
  - Hardware implementation, 96, 209
  - HealthLog daemon, 245, 249
  - HealthLog/stresslog/predictor daemons, 246
  - Heating, Ventilation, and Air conditioning Control (HVAC) techniques, 216
  - Heterogeneous data center, 2, 67, 69, 76
  - Heterogeneous network cluster
    - FPGA mapping file, 21
    - virtual machines, 13, 14, 17, 21, 35, 39, 47, 55, 65, 75, 76, 82, 90, 166, 172, 205
  - Heterogeneous Parallel Hardware (HPA) environments
    - cross-layer programming approach, 4
    - self-adaptive software systems, 4
  - High level synthesis, 76
  - High performance and energy-efficient embedded systems, 194
  - High Performance Computing (HPC) centers, 113
  - High-level languages (C/C++), 61
  - High-Level Synthesis (HLS), 94, 146, 202
  - High-speed communication, 119

HIPPEROS RTOS, 192  
 HLS pragmas, 94  
 HPC applications  
   data locality, 200  
   massive parallelism, 200  
 HPC technology, 163  
 HPE Moonshot, 67  
 HW *Middleware* layer, 12  
 HW virtual machine, 151, 158  
 Hyper-scale resource management, 168

## I

IaaS Scheduler  
   disaggregated architecture of the dReD-  
   Box platform, 47  
   dReDBox, 3, 36, 38, 39, 41, 46, 47, 55  
   Nova Scheduler, 47  
   Openstack's web user interface, 47  
   virtual machine scheduling, 47  
 IBM and Xilinx, 2  
 IBMs Coherent Accelerator Processor Inter-  
   face (CAPI), 68  
 IBM Systems Labs and IBM Research, 90  
 ICT hardware management, 221  
 Implementing data isolation  
   Ethernet layer (L2), 27  
   VLAN level, 27  
 Initialization phase, 254  
 Instruction Set Architecture (ISA), 59, 62  
 Integrated Development Environment  
   (IDE), 131, 192  
 Intel openCL compiler, 68  
 Intel X86\_64 processors, 79, 92  
 Intel Xeon processors, 66  
 Intel XEON vs. Pynq, 102, 104, 105  
 Inter-Process Communication (IPC) proto-  
   cols, 191  
 Interrupt Source Layer (ISL), 68  
 IoT sensor in Victoria, 15  
 IR (Intermediate Representation) level, 123  
 IR-drop analysis, 257  
 Istituto Superiore Mario Boella (ISMB), 59  
 Iterative pre-copy algorithm, 78

## J

Java Native Interface (JNI), 93  
 Java Virtual Machine (JVM), 93  
 Juno PDN simulation framework, 257

## K

Kernel\_accel (gradients/centroids), 99

Key Performance Indicators (KPIs), 131  
 KVM hypervisor, 82

## L

LCI and LCA databases, 65  
 Linux cluster, 10  
 Linux operating system, 77, 121  
 Logistic Regression (LR) training, 94  
 Logistic Regression and K-Means, 4, 88, 89  
 Logistic regression mapreduce, 93  
 Low-cost image processing, 122

## M

Machine learning inference, 90  
   network acceleration, 90  
   SQL query and data analytics, 90  
   storage compression, 90  
   video transcoding, 90, 119  
 Machine learning (MLib), 91  
 Management and host connectivity, 28  
 Many integrated cores, 165  
 MAPE control loop pattern, 144  
 Mapper functions (*gradients\_kernel* and  
   *centroids\_kernel*), 99, 102  
 MapReduce programs, 91  
 Medical X-ray imaging, 5, 183  
 M2DC modular microserver system archi-  
   tecture, 113  
 Memory and hardware accelerators, 3, 35  
 Memory channel virtualization  
   Memory Management Unit (MMU), 26  
   TCP core, 27  
 Memory driver  
   dynamic allocation and deallocation of  
   memory resources, 49  
   hypervisor-level modules, 49  
   NUMA extensions, 49, 51  
   QEMU process, 50  
 Memory hotplug, 49, 50  
 Mesos-GPUs-accelerators, 167  
 Message-Passing Interface (MPI), 10  
 Message Queuing Telemetry Transport  
   (MQTT), 64  
 Micron's automata processor, 150  
 Microservice model, 69, 70, 75, 76  
 Microsoft  
   catapult architecture, 38  
   hypervisor infrastructure, 31  
   network-connected FPGA model, 31  
 Mid-Board Optics (MBO), 44  
 Middleware platform  
   ethernet interface, 20, 157

- for heterogeneous FPGA network clusters, [10, 19](#)
- FPGA application region, [19](#)
- Input Module, [21](#)
- output module, [20, 21](#)
- physical mapping of Kernels, [19](#)
- Migration of FPGA
  - FPGAs in the cloud, [10, 29, 31](#)
  - precopying, [29, 30](#)
- Millions of Instruction Per Second (MIPS), [170](#)
- Misses Per Kilo Instructions (MPKI), [52](#)
- Mixed National Institute of Standards and Technology (MNIST), [100](#)
- Mobile Edge Computing (MEC) paradigm, [64](#)
- Modified memory, [78](#)
- Modular Microserver Data Centre (M2DC) project
  - OpenStack-based middleware, [4, 126](#)
  - System Efficiency Enhancements (SEE), [4, 115, 126](#)
- Molecular dynamics simulations, [10](#)
- Monitoring hardware behavior, [247](#)
- Moore's law, [1, 87](#)
- MPSoCs (Zynq) based on the Pynq platform, [88](#)
- Multi-Root I/O Virtualization (MR-IOV), [120](#)
- Multiple Program Multiple Data (MPMD), [141](#)
- Multiprocessor System-On-Chips (MP-SOCs), [111, 126](#)
- Multi-tenant FPGA hypervisor, [26](#)

**N**

- Netlib Java framework, [93](#)
- Network Address Translation (NAT), [64](#)
- Network Function Virtualization (NFV), [22, 63](#)
- Networking virtualization, [27](#)
- Network Interface Card (NIC), [13](#)
- Network Management Unit (NMU), [28](#)
- Non-uniform Memory Access (NUMA), [50](#)
- NP-hard optimisation problem, [60](#)
- NVIDIA, [118, 119, 135, 144, 145, 163](#)
- NVIDIA GPUs, [135, 144, 145, 163](#)
- NVIDIA Tegra-K1, [190](#)

**O**

- OASIS Topology and Orchestration Specification for Cloud Applications (TOSCA), [69](#)
- Omp Superscalar (OmpSs), [137](#)
- On-Chip Digital Storage Oscilloscope (OCDSO), [257, 259, 260](#)
- Online Photo Services (OPS), [122](#)
- OpenCL or OmpSs/CompS, [135](#)
- OpenCL kernels, [89](#)
- OpenMP standard, [145](#)
- OpenPOWER foundation, [90](#)
- OpenStack cloud orchestration system, [70](#)
- OpenStack FPGA provisioning
  - PCIe device, [14, 15, 17](#)
- OpenStack Nova API service, [157](#)
- OPERA H2020 European project, [83](#)
- OPERA projects
  - deployment of a compact data center on a truck, [3](#)
  - deployment of a virtual desktop infrastructure, [3](#)
  - energy efficiency of cloud infrastructures, [3, 60, 62–64, 164](#)
  - road traffic monitoring application, [3](#)
- Orthoimages, [82](#)
- Overlays, [92](#)

**P**

- Page Miss Handler (PMH), [79](#)
- Partial Reconfiguration (PR), [25](#)
- Particle Swarm Optimization (PSO), [61](#)
- PCI Express, [69, 119](#)
- PCIe-based flash memory
  - disaggregation of, [37, 38](#)
  - inter-rack and intra-rack communication, [38](#)
  - memory blade, [37](#)
  - Xen hypervisor, [37](#)
- PDN frequency domain simulation, [258](#)
- Platform as a Service (PaaS), [113, 124](#)
- Pochoir stencil compiler, [264](#)
- Post-copy and pre-copy migration techniques, [77](#)
- Power delivery network (PDN), [258](#)
- Power service layer, [68](#)
- Poznan supercomputing and networking center, [124](#)
- PREESM and Silexica, [146](#)
- Processor Counter Monitor (PCM), [103](#)
- Processor System (PS), [154](#)
- Programmable Compute Platform (PCP), [151](#)

Programmable Logic (PL), 154  
 Programming model runtime, 131, 136, 137  
 pRouter hosting hardware of Type 2, 175, 177, 178  
 PYNQ open-source framework, 92  
 PYNQ-Z1's ARM cores, 92  
 PySpark, 89  
 Python API, 92, 94, 99

## Q

QoS in Tango, 144  
 Quad Small Form-factor Pluggable (QSFP), 68  
 Quality of Protection (QoP), 131  
 Quality-of-Service (QoS), 164, 244

## R

Radial basis functions  
   Gaussian activation function, 125  
   multilayer perceptrons with sigmoid, 125  
 Radio Access Networks (RANs), 64  
 Radio communication subsystem, 59  
 Radio-frequency (RF) communication interface, 81  
 RAPID EU-funded project, 62  
 RAPL interface, 79  
 Rate monotonic (RM) schedulers, 191  
 Raytrace and data caching workload profiles, 52  
*Razor*, a processor design, 243  
 Real-Time Operating System (RTOS), 5, 183, 192, 195  
*Reconfigurable Acceleration Stack*, 90  
 Redeployment process, 155  
 Renewable energy sources (RES), 217  
 Requirements and design tooling, 131  
 Resilient-Distributed Dataset (RDD), 91  
 Resource and job management systems, 141  
 Resource management and job scheduling, 146  
*Resource Management and Resource Allocation* layers, 12  
 Resource management-OpenStack, 6, 70, 166  
 RISC-based CPU, 53  
 Road traffic monitoring, 3  
 Runtime system, 5, 89, 199, 200, 202, 207, 208, 212

## S

Samsung Exynos5250 SoCs, 119

SAVI FPGA Network Registration, 18  
 SAVI testbed, *see* Smart applications on virtualized infrastructure  
 Self-adaptation manager and monitoring infrastructure, 141  
 Self-organised self-managed (SOSM) framework, 164, 168  
 Service chain scheduler  
   allocation stage, 23  
   connection stage, 23  
 Service Function Chaining (SFC), 22  
 Service-Level Agreements (SLAs) renegotiation, 6, 219, 246  
 Simulated Annealing (SA), 61  
 Single Program Multiple Data (SPMD), 141  
 Single Root I/O Virtualization (SR-IOV), 28  
 SI strategy on CM, 171  
 Slurm controller, 141  
 Smart applications on virtualized infrastructure, 15  
 Software (SW) application and the hardware platform  
   board support package (BSP), 11, 68  
   MPI application, 11, 141  
     *MPI Library*, 10, 11  
     *MPI Program*, 11  
 Software-Defined Memory (SDM) controller, 48  
 Software-Defined Networking (SDN)  
   network connections between CPUs and FPGAs, 14  
   SDN controller, 15, 18, 64  
   network operating system, 15  
 Software/hardware design flow for the cloud, 14, 18  
 Solid development of framework, 252  
 Spark cluster programming framework  
   Blaze, 89  
 Spark executor JVM process, 100  
 Spark MLlib, 91, 92  
 SPECpower benchmark, 65  
 SPynq  
   MPSoC-based platforms, 3  
   PCIe interfaces, 3  
   Xilinx Pynq platform, 3  
   Zynq and heterogeneous systems, 3  
 Standalone manager, 92  
 Static allocation strategy, 71, 74  
 STHEM, *see* Supporting uTilities for Heterogeneous EMbedded image processing  
 StressLog daemon, 141, 208, 249

Suitability Index (SI), [169](#), [170](#), [173](#), [176](#), [179](#)  
 SuperVessel cloud framework, [90](#)  
 Supporting utilities for Heterogeneous Embedded image processing, [192](#), [193](#), [195](#)  
 Support vector machines, [94](#)  
 System Efficiency Enhancements (SEE), [111](#), [115](#), [120](#), [126](#)  
 System-defined infrastructure (SDI), [218](#)  
 System Efficiency Enhancements (SEE), [4](#)  
 System-on-Module (SoM), [195](#)

## T

Tango framework, [141](#)  
 TCP and OpenStack processing, [159](#)  
 Top of the Rack (ToR), [118](#)  
 TOSCA description file, [73](#)  
 Transparent heterogeneous hardware Architecture deployment for eNergy Gain in Operation (TANGO), [130](#)  
 TULIPP reference platform, [5](#), [188](#), [194](#)

## U

Ultra-Low Power (ULP), [59](#), [62](#), [80](#), [81](#), [83](#)  
 UNILogic (Unified Logic), [201](#), [212](#)  
 UNIMEM terminology, [204](#)  
 UniServer approach, The, [6](#), [244](#)  
 University of Toronto, The, [3](#), [10](#), [31](#)  
 Unmanned Aerial Vehicles (UAVs), [5](#), [181](#), [186](#)  
 Urban medium-sized urban DCs, [233](#)  
 Urban micro-DCs, [228](#), [229](#)

## V

VHDL/Verilog, [61](#), [135](#), [145](#)

Virtual and augmented reality (VR/AR), [110](#)  
 Virtual Desktop Infrastructure (VDI), [82](#)  
 Virtualization, [13](#), [15](#), [22](#), [25](#), [28](#), [58](#), [59](#), [63](#), [82](#), [224](#), [226](#), [246](#), [250](#)  
 Virtualized FPGA Resources (VFRs), [151](#)  
 Virtualized Network Functions (VNF), [22](#)  
 Virtual Machines (VMs), [13](#), [14](#), [17](#), [21](#), [31](#), [38](#), [46](#), [47](#), [55](#), [58](#), [61](#), [62](#), [65](#), [75](#), [76](#), [82](#), [90](#), [166](#), [172](#), [205](#), [216](#), [246](#), [250](#)  
 Virtual Rack Managers (vRMs), [167](#), [168](#)  
 Virtualization, [115](#), [120](#)  
 VNF middleware platform, [23](#)  
 Voltage/frequency/refresh rate (V-F-R) point, [245](#)

## W

Warehouse Scale Computer (WSC), [165](#)

## X

X86 64 architecture, [76](#), [92](#)  
 Xeon processors, [10](#), [66](#)  
 X-Gene 2 platform, [251](#)  
 Xilinx dynamic partial reconfiguration, [192](#)  
 Xilinx FPGA accelerators, [90](#)  
 Xilinx memory-mapped IP cores, [16](#)  
 Xilinx Vivado High-Level Synthesis (HLS) tool, [94](#)  
 Xilinx ZC706 development board, [151](#), [157](#)  
 Xilinx Zynq AP SoCs, [92](#)

## Z

ZC702 Evaluation board, [103](#)  
 Zynq FPGA, [97](#), [98](#), [101](#)  
 Zynq xc7z045 FPGA SoC device, [157](#)